SYMPOSIUM OP'24 ADVANCING OCEAN PREDICTION SCIENCE FOR SOCIETAL BENEFITS

SASIP

Down the rabbit hole: the ensemble Kalman filter in the latent spaces of a double variational autoencoder

Ivo Pasmans | Alberto Carrassi | Yumeng Chen | Tobias Finn | Marc Bocquet

"From the moment I fell down that rabbit hole I've been told where I must go and who I must be. I've been shrunk, stretched, scratched, and stuffed into a teapot"
C.S. Lewis, Alice's Adventures in Wonderland, 1865

Introduction

- **Data assimilation (DA) combines a previous guess with observations to provide an estimate for the truth.**
- **The Ensemble Transform Kalman filter[1] (ETKF) is a popular DA method, but not suited for our new *neXtSIM$_{DG}$* sea ice model[2] as the dynamics of the model are nonlinear, some of its variables bounded and its errors are non-Gaussian.**
- **The variational autoencoder (VAE) can transform any distribution to a Gaussian in its latent space.**
- **In this work we apply the ETKF in the latent space of a VAE to overcome the limitations of the ETKF and introduce our solution to how to deal with observations in the latent space.**

Variational Autoencoder

The VAE[3] tries to find parameters θ, ϕ such that an arbitrary and given probability $p_{\mathcal{X}}(\mathbf{x})$ can be related to a standard distribution $\mathcal{N}(\mathbf{z};\mathbf{0},\mathbf{I})$ via

- **Encoder**: $\mathcal{N}(\mathbf{z};\mathbf{0},\mathbf{I}) = \int q_\phi(\mathbf{z}|\mathbf{x}) p_{\mathcal{X}}(\mathbf{x})\,\mathbf{dx}$
- **Decoder**: $p_{\mathcal{X}}(\mathbf{x}) = \int p_\theta(\mathbf{x}|\mathbf{z}) \mathcal{N}(\mathbf{z};\mathbf{0},\mathbf{I})\,\mathbf{dz}$

by maximising the expectation value of the evidence lower bound (ELBO)

$$\mathcal{L}(\theta,\phi,\mathbf{x}) = \int q_\phi(\mathbf{z}|\mathbf{x}) \log p_\theta(\mathbf{x}|\mathbf{z})\,\mathbf{dz} - \int q_\phi(\mathbf{z}|\mathbf{x}) \log \frac{q_\phi(\mathbf{z}|\mathbf{x})}{\mathcal{N}(\mathbf{z};\mathbf{0},\mathbf{I})}\,\mathbf{dz}$$

which, under the assumption that $p_\theta(\mathbf{x}|\mathbf{z}), q_\phi(\mathbf{z}|\mathbf{x})$ are Gaussian, can be found by minimising

$$\frac{1}{2}||\mathbf{x}^{(m)} - \mu_\theta(\mathbf{z}^{(m)})||^2_{\mathbf{\Sigma}_\theta^{-1}(\mathbf{z}^{(m)})} + \log\det \mathbf{\Sigma}_\theta(\mathbf{z}^{(m)}) + D + \log\det \mathbf{\Sigma}_\phi(\mathbf{x}^{(m)}) - \mathrm{Tr}\mathbf{\Sigma}_\phi(\mathbf{x}^{(m)}) - ||\mu_\phi(\mathbf{x}^{(m)})||^2$$

with $\mathbf{x}^{(m)}$ drawn from $p_{\mathcal{X}}(\mathbf{x})$ and $\mathbf{z}^{(m)}$ a point in the so-called latent space.

Single and double VAE-DA

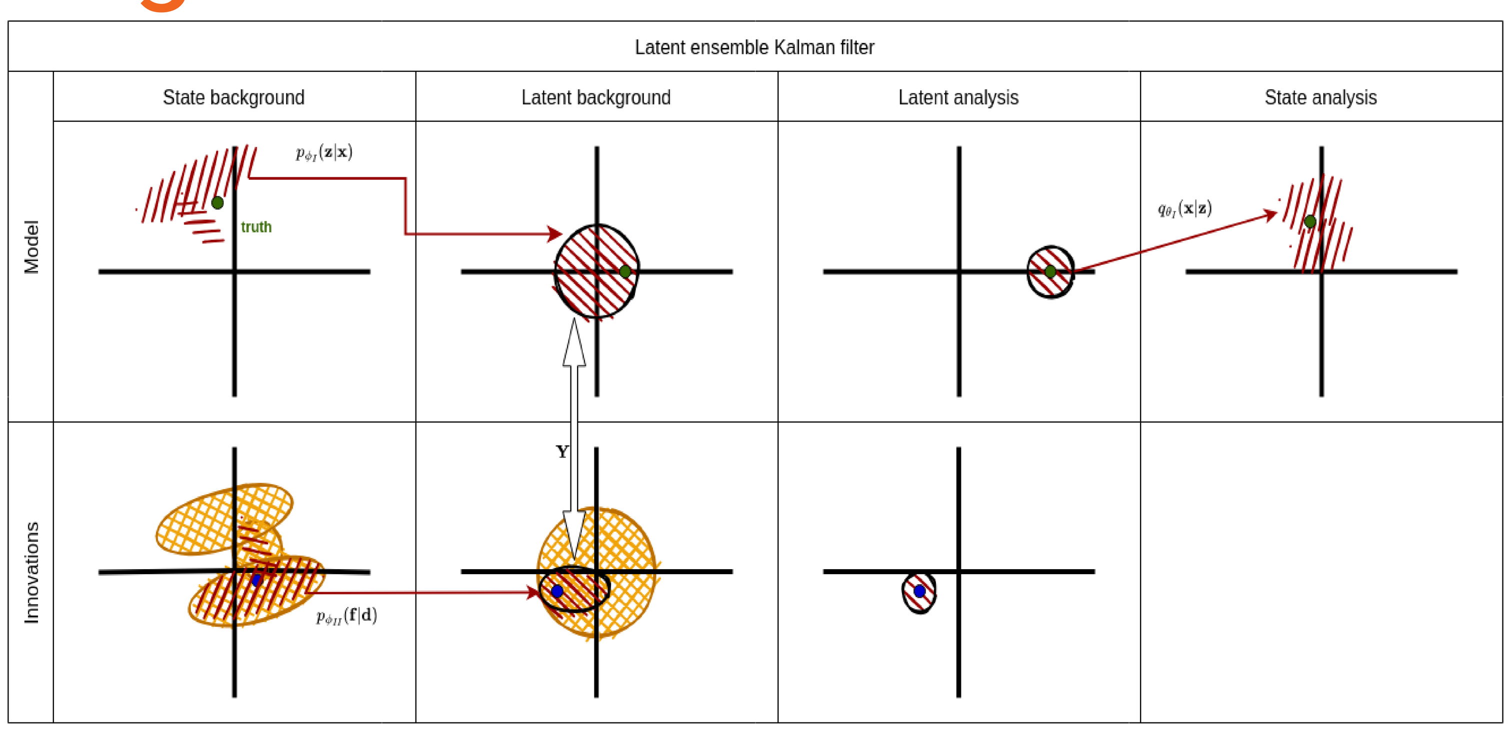

Figure 1: schematic overview of (top) single VAE-DA and (top and bottom) double VAE-DA approach.

ETKF: $\mathbf{X}^{\mathrm{a}} = \mathbf{X}^{\mathrm{f}} + \frac{1}{N-1}\mathbf{X}'^{\mathrm{f}}\mathbf{Y}'^{\dagger}\mathbf{C}^{-1}\overline{\mathbf{d}}\mathbf{1}^{\dagger} + \mathbf{X}'^{\mathrm{f}}(\mathbf{I}-\mathbf{Y}'^{\dagger}\mathbf{C}^{-1}\mathbf{Y}')^{\frac{1}{2}}$ with $\mathbf{C} = \frac{1}{N-1}\mathbf{Y}'\mathbf{Y}'^{\dagger} + \mathbf{R}$

Here the columns of $\mathbf{X} \in \mathbb{R}^{D\times N}$ are the ensemble members, $\mathbf{R}$ the observational error covariance, $\mathbf{Y} = \mathbf{HX}$ the samples taken from the ensemble members, ' indicates that the ensemble mean is removed from the columns and $\overline{\mathbf{d}} = \mathbf{y} - N^{-1}\mathbf{Y1}$, the innovation, i.e. the difference between observations and ensemble prediction for them.
Single-VAE (1st row fig. 1): 1) for each ensemble member use the encoder to draw a latent ensemble member; 2) update the latent ensemble using ETKF; 3) for each latent ensemble member draw a physical ensemble member using the decoder.
Double-VAE (1st and 2nd row fig. 1): 1) Train a 2nd VAE using synthetic innovations created from the ensemble; 2) Add synthetic observational errors and apply encoder to estimate $\mathbf{C}$ in the 2nd latent space; 3) Apply the encoder to innovations without additional observational errors to estimate $\mathbf{Y}, \overline{\mathbf{d}}$ in the latent space.

DA configuration	VAE transformations	1st VAE training	2nd VAE training
No DA	None		
ETKF	None		
Single-clima	States	Offline	
Double-clima	States, innovations	Offline	Online
Single-transfer	State	Online	
Double-transfer	States, innovations	Online	Online

Table 1: different DA configurations used in the experiments.

Model

The model rotates a particle around a circle with the rotation angle depending on polar angle of the particle. Optionally, the radius of the circle may vary in time.

$$\frac{d\psi}{dt} = \alpha_\psi \psi(t) \text{ with } 0 \leq \psi < 2\pi$$
$$\frac{dr}{dt} = \alpha_r \frac{2\pi}{50}\cos\left(\frac{2\pi t}{50}\right)$$

Experiments

Twin-experiments are carried out in which a T=10,000 long model run is used to train the clima VAE. A different T=500 run is used to generate an artificial 'truth'. Every 10 steps an observation of the x-coordinate is generated from this 'truth' with $\mathbf{R} = 0.01$. These observations are assimilated into a 64-member ensemble using the different DA configurations in Table 1. This experiment is repeated 48 times. Performance is measured by the continuous rank probability score[4] (CRPS).

$$CRPS = \mathbb{E}\left[\int_{-\infty}^{\infty} CDF(w) - \mathcal{H}(w - w_{truth})\,\mathrm{d}w\right]$$

- **Experiment I: stationary climatology**. The truth stays on the unit circle ($\alpha_r = 0$). Use of the VAE-DA improves the forecast performance as shown by the lower CRPS values (figure 3). This is especially true for the radius, as the decoder ensures that the analysis members stay on the unit circle (figure 4).

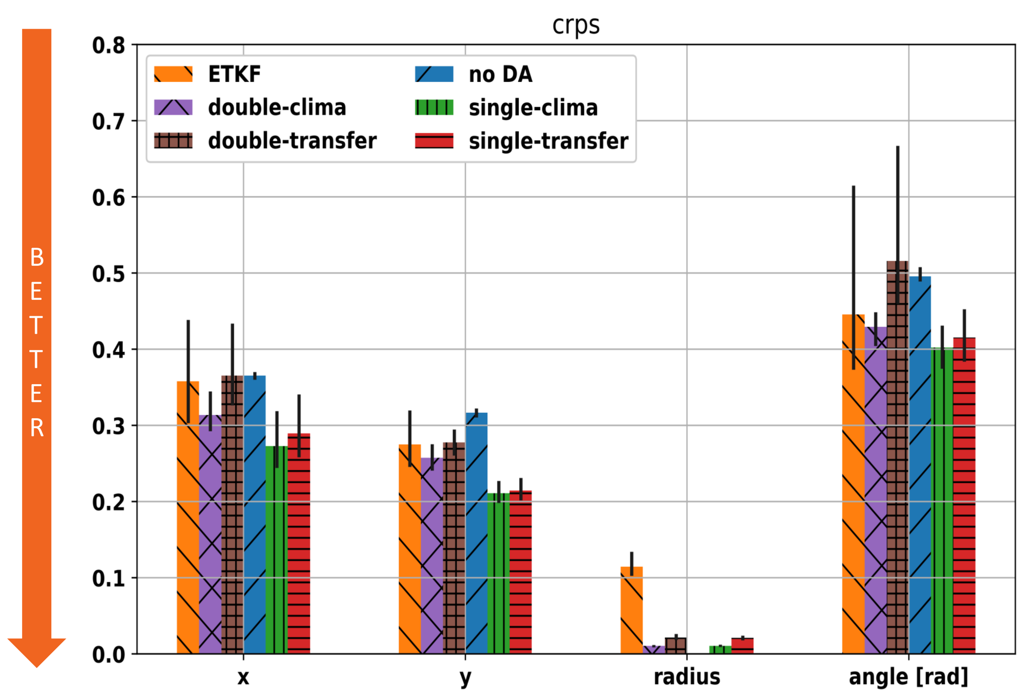

Figure 3: forecast (prior DA) CRPS for different variables and different DA configurations.

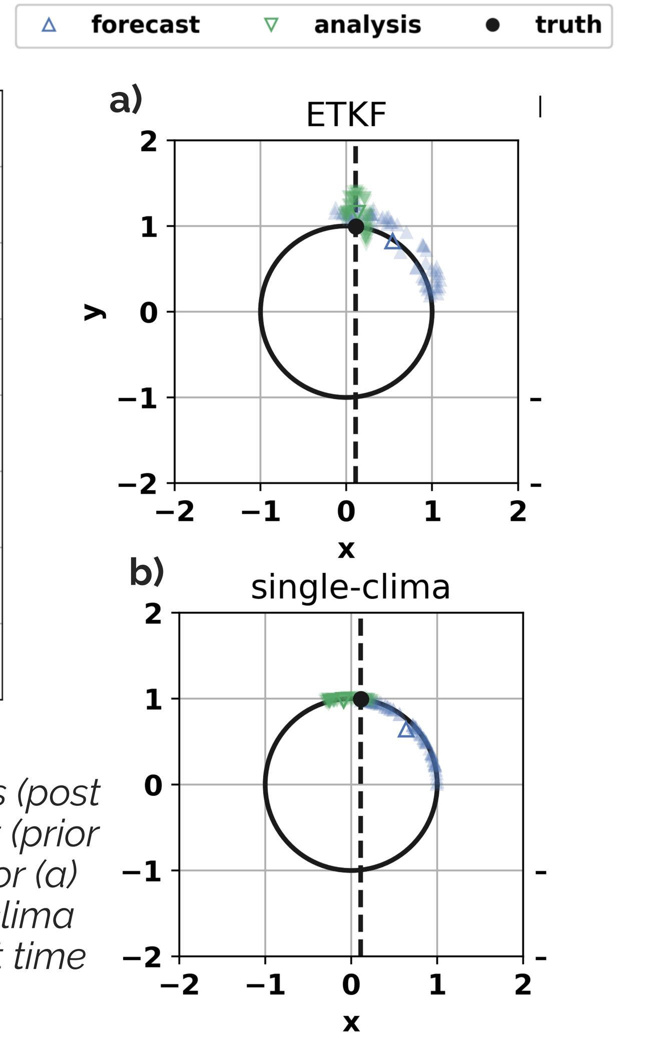

Figure 4: analysis (post DA) and forecast (prior DA) ensembles for (a) ETKF (b) single-clima configurations at time 350.

- **Experiment II: variable climatology.** *clima* configurations fail to capture the variability in the radius of the circle: online training helps with this (figure 5). However, if variability becomes too large even this will fail (figure 6).

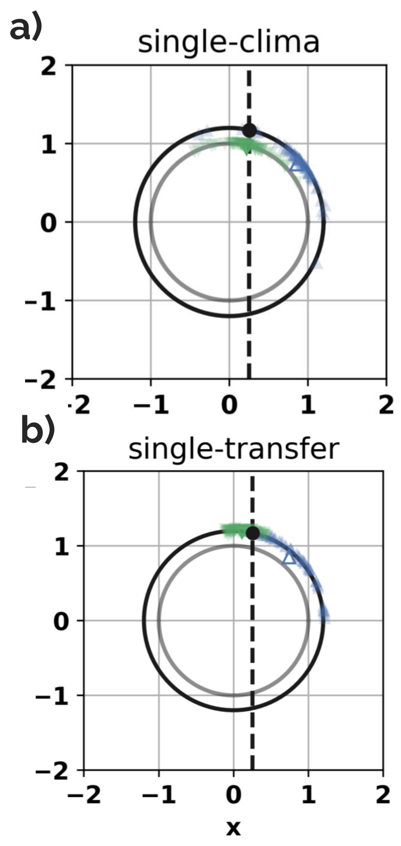

Figure 5 (left): analysis (post DA) and forecast (prior DA) ensembles for (a) single-clima (b) single-transfer configurations at time 350.

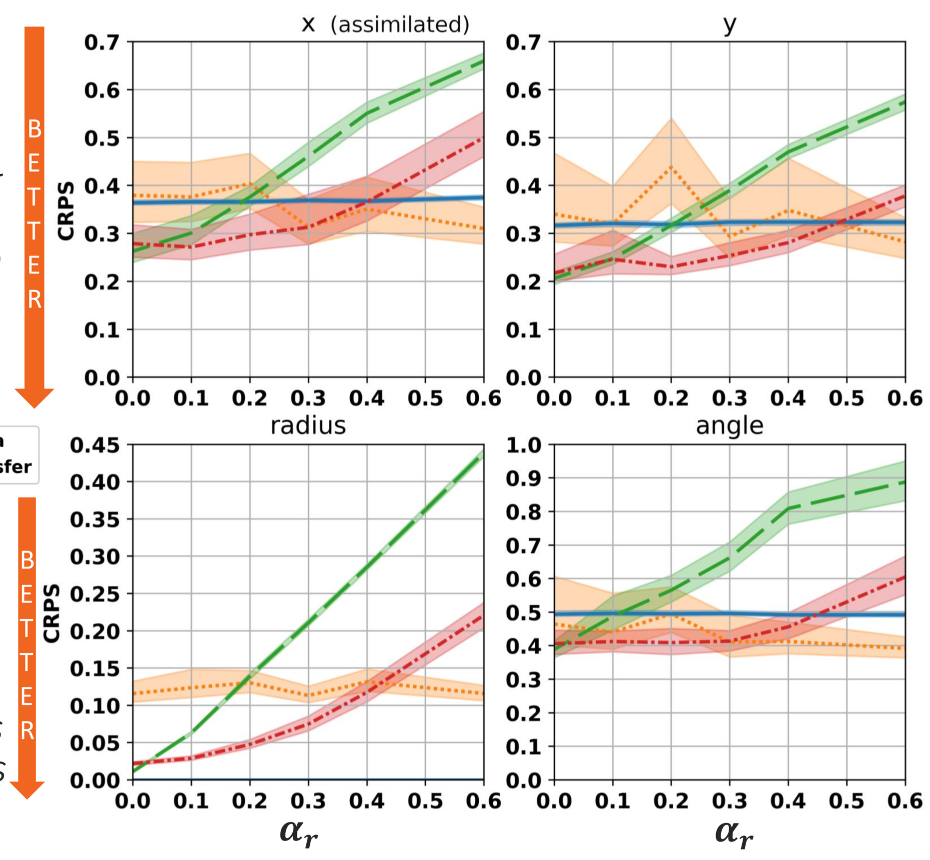

Figure 6 (right): forecast CRPS for different variables and different DA configurations as function of radius change rate.

- **Experiment III: non-Gaussian errors**. Rerun of experiment I, but now the observational errors are réalisations of a skewed-normal distribution (figure 7). The double VAE-DA configuration can deal with skewed, and biased, errors while other configurations cannot (figure 8). Correct error statistics have to be available though.

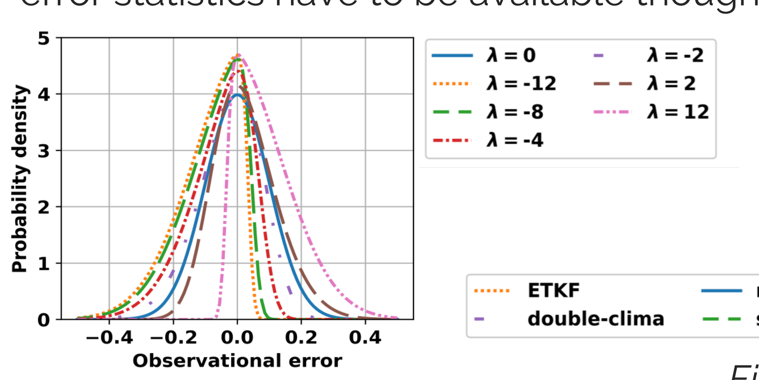

Figure 7: skewed-normal error distribution for different values skewness parameter.

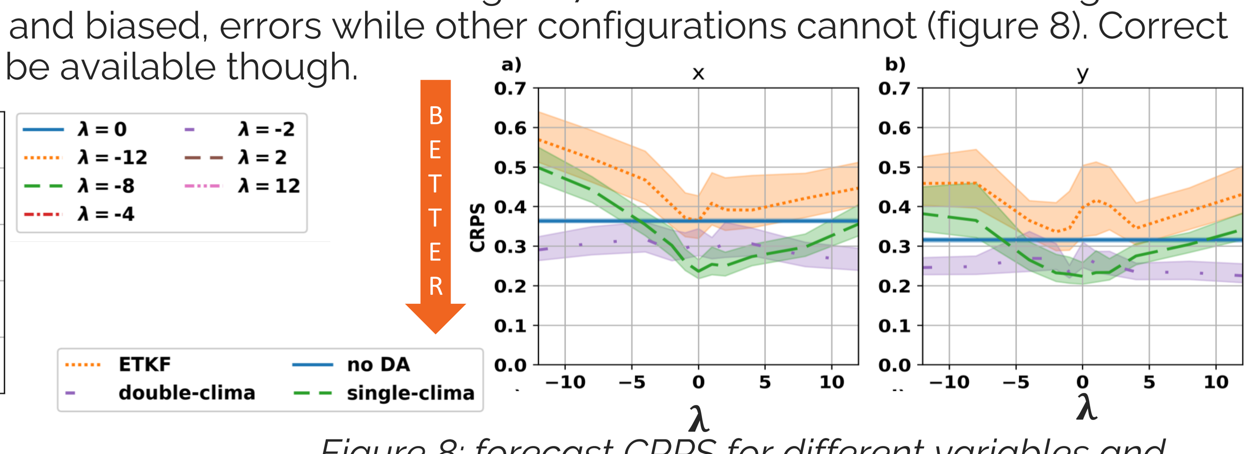

Figure 8: forecast CRPS for different variables and different DA configurations as function of the skewness parameter.

Conclusions

- **Main benefit VAE-DA is that it constraints ensemble members to the physical manifold (in this case a circle).**
- **If that manifold is nonstationary, online training is essential.**
- **Double VAE-DA has the benefit that it averts DA performance degradation when observational errors are biased.**

References:
1. Bishop, C. H., Etherton, B. J., & Majumdar, S. J. (2001). Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Monthly weather review*, *129*(3), 420-436.
2. https://sasip-climate.github.io/
3. Kingma, D. P., & Welling, M. (2019). An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, *12*(4), 307-392.
4. Hersbach, H. (2000). Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, *15*(5), 559-570.

Contact information:
- Department of Meteorology, University of Reading, Whiteknights, RG6 6AH
- Email: **i.c.pasmans@reading.ac.uk**
- https://uk.linkedin.com/in/ivopasmans